

Amendments to the Claims

This listing of claims will replace all prior versions and listings of claims in the above-identified application.

Listing of Claims

1. (Currently Amended) A method comprising:
identifying a plurality of secondary nodes ~~to which an update to data is sent~~, wherein
the identifying comprises sending an update to the plurality of secondary nodes, and
at least one secondary node of the plurality of secondary nodes inserts the update in a respective log of updates, and each of the respective log of updates corresponds to a respective copy of the data; and
sending a notification to each of the plurality of secondary nodes when once all of the plurality of secondary nodes have acknowledged the update.

2. (Previously Presented) The method of claim 1 wherein
each secondary node of the at least one secondary node clears the update from the respective log of updates in response to receiving the notification.

3. (Previously Presented) The method of claim 2 wherein
clearing the update from the respective log comprises updating a start-of-log pointer in the respective log.

4. (Previously Presented) The method of claim 2 wherein
the clearing the update from the respective log comprises updating a pointer to a location in the respective log, wherein
the pointer points to the location if the location contains a next update to clear.

5. (Previously Presented) The method of claim 1 further comprising:
determining that a location of a next update in a first respective log of updates to a first
respective copy of the data at a first secondary node of the secondary nodes
differs from a corresponding location of the next update in a second respective log
of updates to a second respective copy of the data at a second secondary node of
the secondary nodes; and
identifying a set of updates in the first respective log, wherein
each update of the set of updates is not in the second respective log; and
synchronizing the first respective copy and the second respective copy by applying the
set of updates to the second respective copy.
6. (Currently Amended) The method of claim [[1]] ~~5~~ wherein
the determining occurs when a primary node maintaining the data fails.
7. (Previously Presented) The method of claim 1 further comprising:
setting a sent indicator for the update for one of the plurality of secondary nodes when the
update is sent to the one secondary node.
8. (Previously Presented) The method of claim 7 further comprising:
setting a received indicator for the update for the one secondary node when an
acknowledgement of the update is received from the one secondary node.
9. (Previously Presented) The method of claim 8 wherein
the sending the notification to each of the plurality of secondary nodes comprises
determining that a respective sent indicator and a respective received indicator for
the update are set for each of the plurality of secondary nodes.

10. (Currently Amended) A system comprising:
identifying means for identifying a plurality of secondary nodes ~~to which an update to~~
~~data is sent~~, wherein
said identifying comprises sending an update to said plurality of secondary
nodes, and
at least one secondary node of the plurality of secondary nodes inserts the update
in a respective log of updates, and each of the respective log of updates
corresponds to a respective copy of the data[();], and
sending means for sending a notification to each of the plurality of secondary nodes
when once all of the plurality of secondary nodes have acknowledged the update.

11. (Previously Presented) The system of claim 10 further comprising:
clearing means for clearing the update from the respective log of updates in response to
receiving the notification.

12. (Previously Presented) The system of claim 10 further comprising:
determining means for determining that a location of a next update in a first respective
log of updates to a first respective copy of the data at a first secondary node of the
secondary nodes differs from a corresponding location of the next update in a
second respective log of updates to a second respective copy of the data at a
second secondary node of the secondary nodes; and
second identifying means for identifying a set of updates in the first respective log,
wherein
each update of the set of updates is not in the second respective log; and
synchronizing means for synchronizing the first respective copy and the second
respective copy by applying the set of updates to the second respective copy.

13-17. (Cancelled)

18. (Currently Amended) A computer-readable storage medium having a plurality of instructions embodied therein, wherein said plurality of instructions are executable for comprising:

identifying instructions to identify a plurality of secondary nodes ~~to which an update to data is sent~~, wherein

said identifying comprises sending an update to said plurality of secondary nodes, and

at least one secondary node of the plurality of secondary nodes inserts the update in a respective log of updates , and each of the respective log of updates corresponds to a respective copy of the data[;], and

sending instructions to send a notification to each of the plurality of secondary nodes ~~when once~~ all of the plurality of secondary nodes have acknowledged the update.

19. (Currently Amended) The computer-readable storage medium of claim 18 further comprising:

clearing instructions to clear the update from the respective log of updates in response to receiving the notification.

20. (Currently Amended) The computer-readable storage medium of claim 19 wherein the clearing instructions further comprise updating instructions to update a start-of-log pointer in the respective log.

21. (Currently Amended) The computer-readable storage medium of claim 19 wherein the clearing instructions further comprise updating instructions to update a pointer to a location in the respective log, wherein the pointer points to the location if the location contains a next update to clear.

22. (Currently Amended) The computer-readable storage medium of claim 18 further comprising:

determining instructions to determine that a location of a next update in a first respective log of updates to a first respective copy of the data at a first secondary node of the secondary nodes differs from a corresponding location of the next update in a second respective log of updates to a second respective copy of the data at a second secondary node of the secondary nodes; and

second identifying instructions to identify a set of updates in the first respective log, wherein

each update of the set of updates is not in the second respective log; and

synchronizing instructions to synchronize the first respective copy and the second

respective copy by applying the set of updates to the second respective copy.

23. (Previously Presented) A computer system comprising:

a processor for executing instructions, and

a memory to store the instructions, wherein the instructions comprise

identifying instructions to identify a plurality of secondary nodes to which an update to data is sent, wherein

at least one secondary node of the plurality of secondary nodes inserts the update

in a respective log of updates to a respective copy of the data; and

sending instructions to send a notification to each of the plurality of secondary nodes

when all of the plurality of secondary nodes have acknowledged the update.

24. (Previously Presented) The computer system of claim 23 wherein

the instructions further comprise:

clearing instructions to clear the update from the respective log of updates in response to receiving the notification.

25. (Previously Presented) The computer system of claim 23 wherein the instructions further comprise determining instructions to determine that a location of a next update in a first respective log of updates to a first respective copy of the data at a first secondary node of the secondary nodes differs from a corresponding location of the next update in a second respective log of updates to a second respective copy of the data at a second secondary node of the secondary nodes; and second identifying instructions to identify a set of updates in the first respective log, wherein each update of the set of updates is not in the second respective log; and synchronizing instructions to synchronize the first respective copy and the second respective copy by applying the set of updates to the second respective copy.